

Strength Lies in Differences! Improving Strategy Planning for Non-collaborative Dialogues via Diversified User Simulation

Tong Zhang^{♠♥}, Chen Huang^{♠♥}, Yang Deng[◇], Hongru Liang^{♠♥},
Jia Liu[♣], Zujie Wen[♣], Wenqiang Lei^{♠♥*}, Tat-Seng Chua^{*}

♠ Sichuan University ◇ Singapore Management University

♣ Ant Group, China ★ National University of Singapore

♥ Engineering Research Center of Machine Learning and Industry Intelligence,
Ministry of Education, China

{scu.zhangtong, huangc.scu}@gmail.com {lianghongru, wenqianglei}@scu.edu.cn
{jianiulj, zujie.wzj}@antgroup.com ydeng@smu.edu.sg chuats@comp.nus.edu.sg

Abstract

We investigate non-collaborative dialogue agents, which are expected to engage in strategic conversations with diverse users, for securing a mutual agreement that leans favorably towards the system’s objectives. This poses two main challenges for existing dialogue agents: 1) The inability to integrate user-specific characteristics into the strategic planning, and 2) The difficulty of training strategic planners that can be generalized to diverse users. To address these challenges, we propose TRIP to enhance the capability in tailored strategic planning, incorporating a user-aware strategic planning module and a population-based training paradigm. Through experiments on benchmark non-collaborative dialogue tasks, we demonstrate the effectiveness of TRIP in catering to diverse users.

1 Introduction

Non-collaborative dialogues, such as negotiation (He et al., 2018) and persuasion (Wang et al., 2019), occur when the agent and user hold conflicting interests (Deng et al., 2023a,b; Lei et al., 2022). Typically, both parties need to employ various strategies to achieve an agreement favorable to themselves (Keizer et al., 2017; Zhan et al., 2024). As user resistance varies depending on the agent’s strategies (Shi et al., 2019; Dutt et al., 2021), **it is imperative for the agent to perform strategic planning tailored to diverse users**. Relying on a one-size-fits-all strategy can leave the agent vulnerable to others taking advantage due to its lack of adaptability and flexibility (Yang et al., 2021; Deng et al., 2024; Xu et al., 2023).

Recent efforts have resorted to large language models (LLMs) as dialogue agents to perform non-collaborative tasks (Deng et al., 2023d; Fu et al.,

2023; Zhang et al., 2023a). They aim to guide the response of LLMs through mixed-initiative prompts (Chen et al., 2023; Deng et al., 2023d; Zhang et al., 2023a) or incorporating an external strategy planner (Yu et al., 2023; Deng et al., 2023e). However, these initiatives has been criticized regarding its performance in real-world scenarios (Deng et al., 2023e; Kwon et al., 2024), where users have various non-collaborative strategies. We attribute this outcome to the neglect of two crucial aspects: 1) **Existing methods fail to incorporate explicit user-specific characteristics into their strategic planning**, instead relying solely on the conversational history. Importantly, by creating informative representations of individual users, agents can adapt their behaviors and devise tailored strategies (Jang et al., 2020; Yang et al., 2021). 2) **Their training paradigm fails to generate strategic planners that generalize well to diverse users**. Their paradigms are oversimplified, relying on a single user simulator for interactive training. This simulator is restricted in generating varied non-collaborative behaviors, often exhibiting a focus on prioritizing user contentment (Zhang et al., 2023c; Durmus et al., 2023; Bianchi et al., 2024). Essentially, agents trained in this manner are accustomed to engage with a single user exclusively, leading to rigidity and obstinacy when encountering new users with different interaction behaviors (Wang et al., 2023; Safdari et al., 2023).

To provide more evidence for the above analysis, we establish an evaluation protocol, which situates diverse user simulators with varying non-collaborative behaviors. We investigate the limitations of current LLM-based dialogue agents on strategic planning (cf. Section 3 for details). The evaluation results clearly demonstrate that existing agents struggle to tailor their strategies for diverse users, leading to sub-optimal performances.

* Corresponding author.

This limitation compromises the practical utility of these agents, both in functioning as a successful agent in conversational AI and in providing social skills training in pedagogy. **The key challenges lie in making dialogue agents aware of diverse non-collaborative user behaviors and devising tailored strategies for individual users.**

To tackle these challenges, we design a simple yet effective method, called **TRIP**, to improve LLMs' capability in **T**ailored **s**trategic **P**lanning. TRIP includes a user-aware strategic planning module and a population-based training paradigm. Specifically, the strategic planning module incorporates user-specific characteristics into strategic planning using the Theory-of-Mind (ToM) (Premack and Woodruff, 1978; Wimmer and Perner, 1983). This involves analyzing users' mental states and future possible actions during interactions to understand their interests (Yang et al., 2021; Chawla et al., 2023a). Moreover, instead of relying on a solitary user simulator, our population-based training paradigm promotes the adaptation of the strategic planning module to various users, achieved by training it with more diverse user simulators. Each simulator is equipped with extensive sets of non-collaborative strategies and role-playing personas (Chen et al., 2024). As such, TRIP essentially manipulates the experience of the dialogue agent, enabling it to recognize the importance of tailoring strategies for individual users. Our key contributions are concluded below:

- We emphasize the significance of tailoring strategies for diverse users in non-collaborative dialogues. We verify the inadequacies of current LLM-based dialogue agents in this aspect.
- We propose TRIP to achieve tailored strategic planning, which includes a user-aware strategic planning module and a population-based training paradigm.
- We conduct experiments on benchmark non-collaborative dialogue tasks (i.e., negotiation and persuasion). Our findings suggest that TRIP is proficient in catering to diverse users using tailored strategies, consistently outperforming baselines across different tasks.

2 Related Work

Our research is closely tied to the strategic planning and training paradigms to address the non-

collaborative tasks in the era of LLMs. We provide a literature review and highlight our differences.

Strategic planning for non-collaborative dialogues. Recent researches have introduced various methods based on LLMs to enhance their effectiveness in strategic planning. These methods can be categorized into two types: 1) *Developing stimulus prompts to unleash the potential of LLMs.* (Chen et al., 2023) validate the effectiveness of using mixed-initiative prompts to tackle proactive dialogue challenges. (Deng et al., 2023d) and (Zhang et al., 2023a) encourage LLMs to engage in self-reflection to plan their next actions. (Fu et al., 2023) employ self-play simulations to iteratively refine strategic planning by soliciting feedback from other LLMs. Nonetheless, as highlighted by (Deng et al., 2023e), the effectiveness of these approaches is impeded by non-trainable parameters. 2) *Equipping LLMs with an external strategy planner.* The planner is capable of generating prompts at each turn, providing nuanced, instance-specific guidance and control over LLMs. This could be integrated using methods like Monte Carlo Tree Search (Yu et al., 2023) or a plug-in model (Deng et al., 2023e), which can be fine-tuned for improving the strategic planning capability without affecting the functionalities of LLM-powered dialogue agents. However, these methods still struggle to achieve promising results due to their inability to integrate user-specific characteristics into their strategic planning. Complementary to (Deng et al., 2023e), our work investigates the importance of tailored strategic planning by modeling user-related characteristics explicitly. **Training paradigms for non-collaborative dialogues.** Current training paradigms involve the dialogue agent interacting with a single user simulator to enhance its strategic planning capabilities. In specific, (Chawla et al., 2023b) build a user simulator that mimics human-human dialogue data in a supervised manner, while (Yu et al., 2023; Deng et al., 2023e) resort to a role-playing LLM-based user simulator. However, a single user simulator can only represent the behaviors of one or a type of users, potentially leading to the under-representation of other users' behaviors, as evidenced by (Liu et al., 2023; Shi et al., 2019). Therefore, existing training paradigms fail to produce strategic planners that cater to diverse users with varying behaviors. In this paper, our work investigates the importance of tailored strategic planning by diversifying the user's behaviors using population-based training.

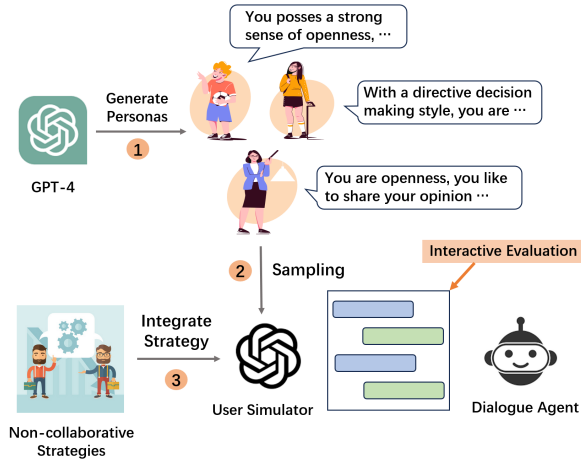


Figure 1: The overall evaluation process.

3 Strategic Planning Evaluation

We introduce a novel evaluation protocol to analyze the limitations of existing LLM-based dialogue agents and highlight their inability to handle users exhibiting various non-collaborative behaviors. The overall evaluation process is illustrated in Figure 1. See more details of our evaluation protocol in Appendix A.

3.1 Evaluation Setup

Evaluation Overview. The environment encompasses various synthetic user simulators showcasing diverse non-collaborative behaviors. In the evaluation process, each dialogue agent must interact with these simulators (Deng et al., 2023e). During their interactions, the dialogue agent and user simulator alternate in employing strategies in their responses with the ultimate aim of maximizing their own self-interest. The interactions continue until the conversational goal is achieved or the maximum number of turns is reached. We gather these interactions and assess the agents’ performances.

Baselines. We consider two representative baselines: *Standard* agent (i.e., vanilla LLM without any modification) and *PPDPP* agent (Deng et al., 2023e), which is current SOTA agent with a trainable external strategy planner¹.

Diverse User Simulators. Our simulators are synthesized with non-collaborative behaviors, guided by their task-relevant personas. As evidenced by previous study (Deng et al., 2023c; Bianchi et al., 2024; Huang et al., 2024), LLMs are limited to demonstrate non-collaborative behaviors. To this

¹Notably, we also consider other existing dialogue agents in our main experiments.

end, we prompt non-collaborative behaviors explicitly into LLMs using the *resisting strategies* that are designed to foil persuasion attempts (Fransen et al., 2015; Tian et al., 2020; Dutt et al., 2021). Initially, we equip LLMs with different personas (Jiang et al., 2023; Zhou et al., 2023b; Zhang et al., 2023b), which are used to select non-collaborative behaviors from the set of *resisting strategies*. Following (Wang et al., 2019; Jiang et al., 2024), we consider two types of personas, including *Big-Five Personality*² (Goldberg, 1992) and *Decision-Making Styles*³ (Scott and Bruce, 1995), together with LLM-generated cohesive description for each fine-grained persona. Additionally, we employ *resisting strategies* outlined by (Dutt et al., 2021) to direct the behavior of simulators. Finally, our mixed-initiative role-play prompt for each agent includes the assigned persona, a set of resisting strategies, and conversation context. These elements aid in guiding user simulators to exhibit diverse non-collaborative behaviors. In total, we develop 300 diverse user simulators for each evaluation task, representing 20 persona categories (i.e., Big-Five Personality \times Decision-Making Styles).

Evaluation Tasks. In line with (Deng et al., 2023d; Wang et al., 2019), we conduct experiments on two benchmark non-collaborative tasks: the price negotiation task, utilizing the test⁴ dataset of Craigslist-Bargain (CB) (He et al., 2018) and the charity persuasion task, employing the test dataset of PersuasionForGood (P4G) (Wang et al., 2019). Notably, the dialogue agents play the role of buyer and persuader, respectively, to accomplish their goals.

Evaluation Metrics. Following (Deng et al., 2023e), we consider three commonly used metrics: Success Rate (SR), Average Turn (AT) and Sale-to-List Ratio (SL%). The SR measures effectiveness by the percentage of goal achievement within a maximum number of turns, while AT measures efficiency by the average number of turns required to achieve the goal. As for the CB task, we additionally adopt the SL% (Zhou et al., 2019) to determine the effectiveness of goal completion. Formally, the SL% is expressed as $(P_{deal} - P_{target}^{seller}) / (P_{target}^{buyer} - P_{target}^{seller})$, where P_{deal} is the final deal price, P_{target}^{buyer} and P_{target}^{seller} are the target prices of both parties. A higher SL% repre-

²Openness, Conscientiousness, Extraversion, Agreeableness, and Neuroticism

³Directive, Conceptual, Analytical, and Behavioral

⁴Our data split follows the previous study (Deng et al., 2023e; Wang et al., 2019).

Personas		Price Negotiation			Persuasion for Good	
		SR \uparrow	AT \downarrow	SL% \uparrow	SR \uparrow	AT \downarrow
Big Five	Openness	0.76 \uparrow 0.23	6.66 \uparrow 0.63	0.34 \uparrow 0.12	0.47 \uparrow 0.34	8.92 \uparrow 1.00
	Conscientiousness	0.69 \uparrow 0.25	7.20 \uparrow 1.04	0.27 \uparrow 0.06	0.39 \uparrow 0.33	8.90 \uparrow 1.10
	Extraversion	0.74 \uparrow 0.16	6.17 \uparrow 1.47	0.39 \uparrow 0.15	0.45 \uparrow 0.35	8.73 \uparrow 1.25
	Agreeableness	0.40 \uparrow 0.01*	6.82 \uparrow 0.71	0.28 \uparrow 0.06	0.18 \uparrow 0.12	9.85 \uparrow 0.13*
	Neuroticism	0.31 \downarrow 0.02*	6.81 \uparrow 1.12	0.20 \downarrow 0.02*	0.12 \uparrow 0.02*	9.78 \uparrow 0.14*
Decision	Analytical	0.37 \uparrow 0.04*	7.07 \uparrow 0.61	0.26 \uparrow 0.06*	0.16 \uparrow 0.09	9.43 \uparrow 0.56*
	Directive	0.41 \uparrow 0.05*	6.71 \uparrow 1.48	0.18 \downarrow 0.03*	0.12 \downarrow 0.02*	9.31 \uparrow 0.62
	Behavioral	0.78 \uparrow 0.25	6.45 \uparrow 1.20	0.39 \uparrow 0.16	0.53 \uparrow 0.37	8.94 \uparrow 1.04
	Conceptual	0.77 \uparrow 0.23	6.62 \uparrow 0.78	0.42 \uparrow 0.17	0.49 \uparrow 0.36	9.02 \uparrow 0.94
Overall Performance		0.58 \uparrow 0.14	6.72 \uparrow 1.01	0.31 \uparrow 0.09	0.32 \uparrow 0.23	9.20 \uparrow 0.76

Table 1: The performance of the *PPDPP* dialogue agent testing across various personas of user simulators. **Red (Blue)** indicates the increased (decreased) performance compared to *Standard* dialogue agent. The symbol \star indicates that this performance exhibits minimal variation, specifically within a 5% range of the maximum value. The effectiveness of *PPDPP* varies significantly across different user personas.

sents the buyer gets more benefits from the deal. If failing to reach a deal at the end, we set SL% as 0.

3.2 Experimental Findings

We analyze the performances of existing dialogue agents across user simulators with various non-collaborative behaviors. Specifically, we assess the advancements of *PPDPP* compared to the *Standard* agent. As illustrated in Table 1, while *PPDPP* shows a notable improvement in overall performance, it does not adapt well to users employing different non-collaborative strategies. Its effectiveness varies significantly among users with different personas, with its advantage over the *Standard* not being significant in 17.77% of cases (e.g., it increases SR by 0.02 for *Analytical* in price negotiation.), and even performing worse than the *Standard* in 8.88% of cases (e.g., it decreases SR by 0.02 for *Neuroticism* in price negotiation). This motivates the need for a dialogue agent to perform strategic planning tailored to diverse users⁵.

4 TRIP: Tailored Strategic Planning

To enhance LLMs’ tailored strategic planning, we propose an effective method TRIP, which develops an external planner by modeling user characteristics and training with diverse user simulators. As illustrated in Figure 2, our TRIP includes a user-aware strategic planning module and a population-based training paradigm. The former aims to explicitly model user characteristics (e.g., mental states and future actions), while the latter incorporates diverse user simulators for training simultaneously.

⁵We find that other baselines also have similar issues, as detailed in Section 5.

4.1 User-Aware Strategic Planning

TRIP aims to explicitly infer user characteristics and then incorporate them into the strategic planning module, parameterized by a trainable BERT. In particular, building upon the advanced *Theory-of-Mind* capability of LLMs (Sap et al., 2022; Moghaddam and Honey, 2023), TRIP captures users’ mental states and future possible actions during interactions to understand their interests and predicts how TRIP’s responses may influence them. In this case, mental states pertains to what they aim to accomplish, such as the target price or whether they will donate, while future actions relates to what the user is likely to discuss next (Hu et al., 2023; Zhou et al., 2023a). Formally, given the dialogue history $D = (u_1^{sys}, u_1^{usr}, \dots, u_t^{sys}, u_t^{usr})$, where u_i^{sys} and u_i^{usr} denote the i -th utterances of both parties and t is the number of utterances, we feed the dialogue history D into the LLM and prompt it to infer mental states \mathcal{M} and future actions \mathcal{F} in an open-ended manner, i.e., $P_{LLM}(\mathcal{M}, \mathcal{F}|D)$. Subsequently, we feed the $\{\mathcal{M}, \mathcal{F}, D\}$ into the strategy planner π_θ to predict the next strategy. The output space of π_θ is a set of strategies⁶ pre-defined by (Deng et al., 2023e; Wang et al., 2019), each of them is attached with a pre-defined natural language instructions.

4.2 Population-based Training Paradigm

Given that a single user simulator tends to favor limited behaviors while under-represents others (Shi et al., 2019; Liu et al., 2023), we explore training a dialogue agent using a set of user simulators employing different non-collaborative strategies to accommodate diverse users. To achieve this, we

⁶e.g., the elicitation of specific emotions to influence other.

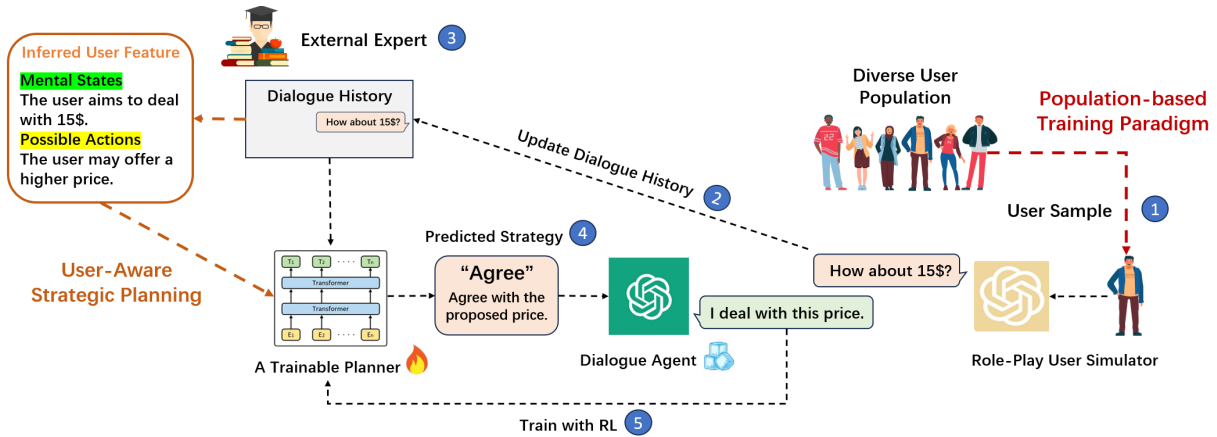


Figure 2: TRIP Overview. This method includes a user-aware strategic planning module (UASP) and a population-based training paradigm (PBTP). The UASP incorporates user-specific characteristics into strategic planning using the Theory-of-Mind (ToM). The PBTP diversifies training user simulators to promote agents’ adaptation. We use numbers to indicate the overall process of TRIP.

propose a population-based reinforcement learning (RL) training paradigm, which aims to enhance the adaptability of a dialogue agent to new user groups by training with larger and more diverse populations (Charakorn et al., 2020). We offer a comprehensive explanation of this approach below.

Population Setup. Similar to Section 3.1, we build 40 diverse user simulators, each embodying a specific persona description. We ensure an balanced representation of each persona category within our user simulators for population-based RL training. We denote these simulators as $K = k_1, k_2, \dots, k_{40}$. During each iteration, we sample among K using a distribution p , allowing the dialogue agent S to interact with it. The distribution p is initialized based on the frequency of various personas.

Reward Design. Following (Deng et al., 2023e), we prompt LLMs to judge the conversation progress at each turn and transform it into scalar rewards. Specifically, in the negotiation task, we employ a separate GPT3.5 (OpenAI, 2022) to assess whether both parties have reached a deal. In the persuasion task, we ask the GPT3.5-based user simulator to express its willingness to donation. Our rewards are determined based on three situations: 1) Successful goal achievement by the dialogue agent results in a significant positive reward, defined as 1.0 in the charity persuasion task and the value of SL% in the price negotiation task. 2) Failure to achieve goals leads to a substantial negative reward of -1.0 for the dialogue agent. 3) Furthermore, we assign a small negative reward (-0.1) per turn to penalize the lengthy conversation, which promotes the efficient goal achievement.

Optimization. During RL training, we maximize the expected reward of the strategy planner π_θ by utilizing the REINFORCE algorithm (Williams, 1992): $\theta \leftarrow \theta - \alpha \nabla \log \pi_\theta R_t$, where θ denotes the trainable parameter of the strategy planner, α denotes the learning rate, and R_t is the total reward accumulating from turn t to the final turn T : $R_t = \sum_{t'=t}^T \gamma^{T-t'} r_{t'}$, where γ is a discount factor.

5 Experiments

This sections aims to evaluate the effectiveness of our TRIP, following the evaluation protocol proposed in Section 3.1. We initially report the overall performances of dialogue agents in Section 5.1. Next, we conduct an in-depth analysis to reveal the tailored strategies of TRIP in Section 5.2. Finally, we perform ablation studies in Section 5.3 to sort out the performance variation of different user awareness and training population, and find a dominant predictor for the tailored strategic planning.

LLM-based baselines. We consider LLM-based dialogue agents with two types of strategic planning modules, as discussed in Section 2. 1) Prompt-based planning, including *Standard*, *Pro-CoT* (Deng et al., 2023d) and *ICL-AIF* (Fu et al., 2023), which use mixed-initiative prompts, CoT, and AI feedback to select next strategies, respectively. 2) External strategy planners, including *GDP-MCTS* (Yu et al., 2023) and *PPDPP* (Deng et al., 2023e), which utilize Monte Carlo Tree Search and a trainable plug-in for determining next-step strategies, respectively. Note that all baselines fail to model user-specific characteristics explicitly and are trained using one user simulator. Imple-

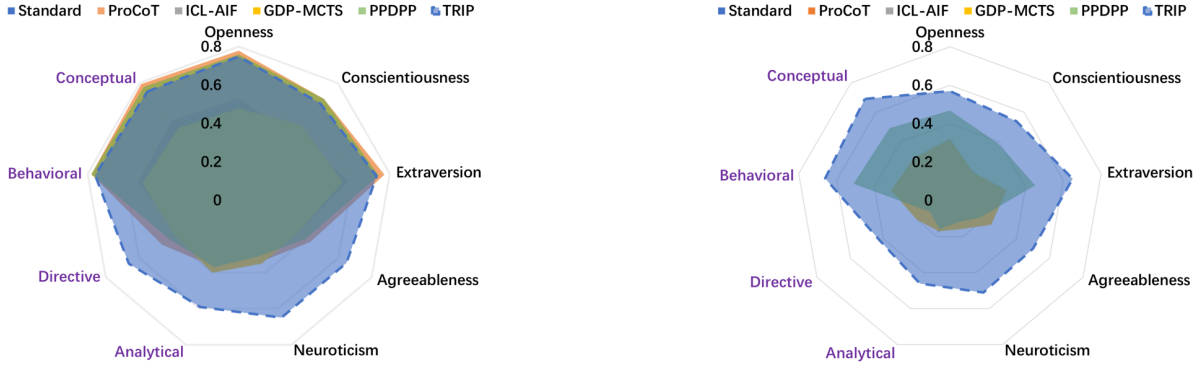


Figure 3: The agents performance across various personas. We report their success rate on two tasks, namely price negotiation (*Left*) and charity persuasion (*Right*). TRIP achieves balanced improvements on all personas, significantly outperforming other agents by a considerable margin. Due to limited space, we report other results using different metrics in Appendix D.

mentation details are presented in Appendix B.

Evaluation Metrics. We use the same automatic metrics mentioned in section 3.1. Furthermore, we conduct human evaluation to assess the practical effectiveness of these dialogue agents. See more details of human evaluation in Appendix C.

5.1 Overall Performance

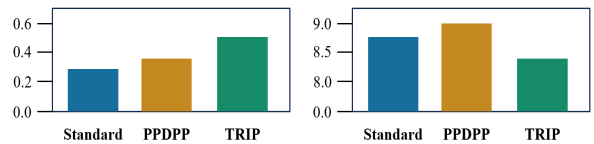
We evaluate the overall and fine-grained performance of all agents using automatic metrics in Table 2 and Figure 3. Additionally, we report human evaluation in Figure 4 to gauge their performance during interactions with real users.

TRIP is a promising method for achieving effective non-collaborative strategies tailored for diverse users. As illustrated in Table 2, TRIP significantly outperforms all the baselines with a noticeable margin across two tasks. It not only efficiently achieves the conversational goal (less AT) but also effectively accomplishes tasks (higher SR and higher SL%). Moreover, as depicted in Figure 3, TRIP shows balanced improvements across different user personas, significantly outperforming other agents by a substantial margin, in contrast to the biased improvements of *PPDPP* in Section 3.2. This suggests that TRIP is capable of generating strategies that generalize well to diverse users. This also implies that the behavior pattern of a single LLM-based user simulator is limited in scope. Moreover, our human evaluation results in Figure 4 show our TRIP largely outperform the *Standard* and *PPDPP* when interacting with real users. Notably, we observed that *PPDPP* does not consistently surpass the *Standard* approach across the two tasks. For instance, while it achieves a higher success rate in the negotiation task, it neces-

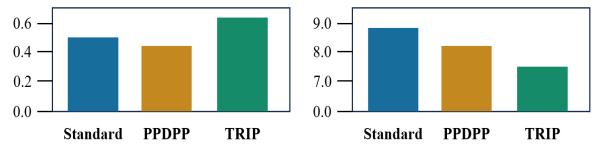
sitates more interaction rounds. This evidences the effectiveness and practical utility of our proposed TRIP.

Agents	Price Negotiation			Persuasion for Good	
	SR \uparrow	AT \downarrow	SL% \uparrow	SR \uparrow	AT \downarrow
Standard	0.4444	7.73	0.2222	0.0930	9.96
ProCoT	0.6040	7.62	0.2307	0.1833	9.90
ICL-AIF	0.3411	8.42	0.2503	0.1667	9.91
GDP-MCTS	0.4444	7.63	0.2401	0.2466	9.74
PPDPP	0.5855	6.72	0.3144	0.3233	9.20
TRIP (<i>Ours</i>)	0.6888	6.34	0.4096	0.5533	8.51

Table 2: Overall evaluation. TRIP is promising for achieving effective non-collaborative strategies.



(a) Success Rate (Left) and Average Turn (Right) on Price Negotiation



(b) Success Rate (Left) and Average Turn (Right) on Charity Persuasion

Figure 4: Human Evaluation Results. TRIP shows a high practical utility to deal with real users.

5.2 Strategy Analysis

In this section, we analyze the effectiveness of our TRIP in tailored strategic planning. Specifically, in each user interaction, we gather the strategies employed by each agent at every turn and combine them in a sequential order to form a strategy sequence. Then, we compare the strategy sequences

User Persona: Openness	User Persona: Neuroticism
(1st Round) User: Hello. How are you? PPDPP: [Task-related Inquiry] Very well. I'm just up organizing info for my charity called Save the Children. Have you heard of this charity before? TRIP: [Task-related Inquiry] I'm doing great, thank you. I wanted to talk to you about a charity called Save the Children. Have you heard of it before?	(1st Round) User: Hello. How are you? PPDPP: [Task-related Inquiry] I wanted to talk to you about a charity called Save the Children. Have you heard of it before? TRIP: [Task-related Inquiry] Have you ever heard of a charity called Save the Children? They work to help fight poverty around the world.
(2nd Round) User: [Neutral] No, I have not. Can you tell me more? PPDPP: [Credibility Appeal] Save the Children works to promote children's rights, provide relief and support to children. TRIP: [Logical Appeal] Yes, Save the Children provides education, healthcare, and other essential services to help children thrive. I think it's a really important cause.	(2nd Round) User: [Counter Argument] No, I have not. That's not my business. There are other people who are richer. PPDPP: [Credibility Appeal] Save the Children works to promote children's rights and protect them from harm. TRIP: [Personal-related Inquiry] How much money do you spend daily on extras like a coffee or treat? Every donation can make a big difference.
(3rd Round) User: [Source Derogation] What percentage of the money goes to the children? My money probably won't go to the right place. PPDPP: [Credibility Appeal] I can assure you that your contribution will be put towards making a real impact in the lives of those who need it most. TRIP: [Emotion Appeal] I understand your concern, and it's important to make sure that you feel comfortable. Save the Children is a highly reputable charity that has been operating for over 100 years.	(3rd Round) User: [Self Pity] I have my own children to support. PPDPP: [Credibility Appeal] I can assure you that Save the Children is committed to being as transparent as possible about their finances. TRIP: [Personal Story] As a parent myself, I strongly believe that every child deserves to experience joy and fulfillment. That's why I've contributed \$100 to support this cause.

Figure 5: Case study on the charity persuasion task (Top-3 conversation rounds). The user resisting strategies and agent strategies are marked in **bleu** and **red** respectively. While *PPDPP* repeats its strategy usage pattern to different user types, *TRIP* effectively tailor its strategies for different users. When dealing with the *Openness* persona (*Left*), *TRIP* introduces the charitable organization and evoke specific emotions to sway users' decision. Conversely, in addressing the *Neuroticism* persona (*Right*), *TRIP* tends to discuss personal experiences related to charity and employs reasoning persuade the user.

Models	Intra-Persona↓	Inter-Persona↑
Standard	24.93	13.51
ProCoT	21.37	15.65
ICL-AIF	22.84	15.33
GDP-MCTS	20.72	16.09
PPDPP	19.37	17.28
TRIP (<i>Ours</i>)	16.14	20.26

Table 3: The strategy distribution of different agents. The Intra-Persona metric donates the average distance for a particular persona. The Inter-Persona metric donate the average distance for different personas. *TRIP* achieves the best performance, showcasing its effectiveness in devising tailored strategies for diverse users.

employed by different agents. We utilize BERT (Devlin et al., 2018) and the t-SNE method (Van der Maaten and Hinton, 2008) to encode each strategy sequence into an embedding vector. Subsequently, we use the Euclidean distance measure to calculate the average distance between any two strategy sequences used by agents with the same persona, as well as the average distance between any two strategy sequences used by agents with different personas. This is akin to the metrics (i.e., the Intra-Class and Inter-Class analysis) used in the metric learning community (Roth et al., 2019) and we term them as the Intra-Persona and Inter-Persona. The results are shown in Table 3.

TRIP demonstrates a greater awareness of population dynamics, resulting in reduced variance across specific user simulators. As shown in Ta-

ble 3, *TRIP* achieves the lowest Intra-Persona and the highest Inter-Persona. This indicates that the strategy sequences of *TRIP* exhibit similarity when interacting with users sharing the same personas and non-collaborative behaviors. Also, these sequences are distinct when compared to users with different personas. This further reveals that *TRIP* holds advantages in devising tailored strategies for diverse users.

For better understanding, we present a case study in Figure 5 and examine the strategy sequence employed by *PPDPP* and *TRIP* in an charity persuasion task. Specifically, *PPDPP* repeats its strategy usage pattern to different user types, briefly using of credentials and citing organizational impacts to establish credibility and earn the persuadee's trust. In contrast, *TRIP* demonstrates a deeper understanding of the users and provides more tailored strategies. When dealing with the *Neuroticism* persona, *TRIP* tends to discuss personal experiences related to charity and employs reasoning persuade the user. Conversely, in addressing the *Openness* persona, *TRIP* introduces the charitable organization and evoke specific emotions to sway users' decision. The strategy sequence used by *TRIP* is believed to be more persuasive, as demonstrated by (Barford and Smillie, 2016; Wang et al., 2019), stating that the *Openness* users are inclined to embrace novelty and be easily influenced by emotions, while the *Neuroticism* users are more likely to be influenced by others' personal experiences. In this regard, we

Models	Price Negotiation			Persuasion for Good	
	SR \uparrow	AT \downarrow	SL% \uparrow	SR \uparrow	AT \downarrow
TRIP	0.6888	6.34	0.4096	0.5533	8.51
TRIP _{w/o UA}	0.6988	6.38	0.3881	0.5133	8.69
TRIP _{w/o POP}	0.5766	7.00	0.3505	0.4400	8.95
TRIP _{w/ 10 POP & w/o UA}	0.6377	6.73	0.3543	0.4700	8.79
TRIP _{w/ 10 POP}	0.6700	6.12	0.3537	0.4733	8.72
PPDPP	0.5855	6.72	0.3144	0.3233	9.20

Table 4: The evaluation results of ablation study. The user-aware strategic planning module and population-based training are effective to improve agents and complement each other.

believe that these strategic differences may provide valuable insights for the future research on the non-collaborative dialogues.

5.3 Ablation Study

This section aims to sort out the performance variation of different user awareness and training population. To analyze the effectiveness of each design, we consider the following variants of TRIP.

- **TRIP_{w/o POP}**: We eliminate the population-based training approach from TRIP and instead have TRIP engage with a single fixed LLM-based user simulator for training, without any specific role-playing persona.
- **TRIP_{w/o UA}**: We remove the user-aware strategic planning module, and only takes the conversation history as inputs to plan next strategies.
- **TRIP_{w/ 10 POP}**: It utilizes 10 personas for population training, each simulator is randomly selected from a pool of 20 persona categories.
- **TRIP_{w/ 10 POP & w/o UA}**: In this variant, we remove the user-aware strategic planning module from TRIP *w/ 10 POP*.

We summarize the overall performance of each model variation Table 4. Based on these results, we draw the following observations:

User-aware strategic planning and population-based training paradigm are both effective to produce tailored strategic planning. Specifically, compared to TRIP_{w/o UA}, we note TRIP improves the persuasion success rate (0.3233 \rightarrow 0.4400) and the deal benefit SL% (0.3144 \rightarrow 0.3505). This suggests that incorporating user mental states and future actions can assist the agent in developing more effective strategies. Notably, this variant slightly decreases the deal success rate (0.6988 \rightarrow 0.6888). This can be attributed to the fact that deeply modeling user characteristics may inadvertently decrease the seller’s willingness to engage in the deal, as the

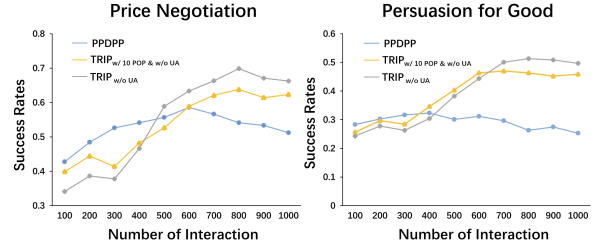


Figure 6: The test performance of different number of training user simulators. *PPDPP* converges easily but has a limited upper bound in terms of performance.

focus is on maximizing one’s own benefits. Moreover, compared to TRIP_{w/o POP}, we observe that TRIP yield positive improvements across all metrics, such as significant increase in SL% (0.3505 \rightarrow 0.4096). This demonstrates that diversifying the behaviors of training user simulators effectively improves the agent’s performance.

Diverse training populations is more beneficial to improve the adaptability of dialogue agents, but it may also present additional training challenges. As shown in Table 4, compared to TRIP_{w/o UA} and TRIP_{w/o POP}, we find that diverse training populations is more important for TRIP’s superiority. Moreover, we find that TRIP_{w/o UA} demonstrates higher performances than TRIP_{w/ 10 POP & w/o UA} and *PPDPP* (i.e., A single fixed user simulator). To provide a detailed understanding of the impact of the number of training user simulators, we present their test performance of in 1000 training interactions, as depicted in Figure 6. Particularly, during the initial 400 interactions, we observe that TRIP_{w/o UA} and TRIP_{w/ 10 POP & w/o UA} exhibit slower convergence compared to *PPDPP*. This suggests that not keeping the training user simulator fixed can introduce instability in the initial training phase, as also noted in (Lewis et al., 2017). However, beyond 500 interactions, the training process of TRIP_{w/o UA} stabilizes, leading to a significant performance enhancement, surpassing the other two agents. Additionally, it is observed that *PPDPP*’s performance declines after specific interactions (e.g., 600 in price negotiation), suggesting that extensive interactions with a single user simulator cannot consistently enhance agents’ performance.

6 Conclusion

In this study, we investigate the inadequacies of current LLM-based dialogue agents in catering in diverse non-cooperative users. To address this, we

propose TRIP, a method designed to tailor strategic planning for non-collaborative dialogues. The idea behind our TRIP is simple, involving a user-aware strategic planning module and a population-based training paradigm. Experimental results across diverse users demonstrate the superior effectiveness and efficiency of TRIP. We consider our work as laying the groundwork for enhancing the adaptability and flexibility of non-cooperative dialogue agents in the era of LLMs. Moving forward, we plan to further explore the potential of population-aware agents in reducing the capital expenditure associated with training and coaching novice agents.

Limitations

In this section, we discuss the limitations of this work from the following perspectives:

Sensitivity of Prompts. Similar to other studies on prompting LLMs (Deng et al., 2023d), the evaluation results are expected to be influenced by the prompts. Following (Deng et al., 2023e), we employ the mixed-initiative format to formulate our prompts, as it offers stability and control. The impact of prompts and their optimality present important areas of investigation within LLMs, calling for exploration in future studies.

Limited Non-collaborative Tasks. We only conduct our experiments on the two non-collaborative dialogue tasks (i.e., price negotiation and charity persuasion) due to their status as classic and widely-recognized benchmarks (Deng et al., 2023d; Chawla et al., 2023a). In the future, we plan to apply our proposed TRIP in a broader range of non-collaborative dialogue scenarios (Zhang et al., 2024; Zhou et al., 2023b; Zhang et al., 2023b).

Acknowledgements

This work was supported in part by the National Natural Science Foundation of China (No. 62272330 and No. 62206191); in part by the Natural Science Foundation of Sichuan (No. 2023NSFSC0473); in part by the Fundamental Research Funds for the Central Universities (No. 2023SCU12089 and No. YJ202219); in part by the Singapore Ministry of Education (MOE) Academic Research Fund (AcRF) Tier 1 grant (No. MSS24C004).

References

Kate A Barford and Luke D Smillie. 2016. Openness and other big five traits in relation to dispositional

mixed emotions. *Personality and individual differences*, 102:118–122.

Federico Bianchi, Patrick John Chia, Mert Yuksekgonul, Jacopo Tagliabue, Dan Jurafsky, and James Zou. 2024. How well can llms negotiate? negotiationarena platform and analysis. *arXiv preprint arXiv:2402.05863*.

Rujikorn Charakorn, Poramate Manoonpong, and Nat Dilokthanakul. 2020. Investigating partner diversification methods in cooperative multi-agent deep reinforcement learning. In *Neural Information Processing: 27th International Conference, ICONIP 2020, Bangkok, Thailand, November 18–22, 2020, Proceedings, Part V 27*, pages 395–402. Springer.

Kushal Chawla, Weiyang Shi, Jingwen Zhang, Gale Lucas, Zhou Yu, and Jonathan Gratch. 2023a. Social influence dialogue systems: A survey of datasets and models for social influence tasks. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 750–766.

Kushal Chawla, Ian Wu, Yu Rong, Gale Lucas, and Jonathan Gratch. 2023b. *Be selfish, but wisely: Investigating the impact of agent personality in mixed-motive human-agent interactions*. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 13078–13092, Singapore. Association for Computational Linguistics.

Maximillian Chen, Xiao Yu, Weiyang Shi, Urvi Awasthi, and Zhou Yu. 2023. *Controllable mixed-initiative dialogue generation through prompting*. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 951–966, Toronto, Canada. Association for Computational Linguistics.

Nuo Chen, Yan Wang, Yang Deng, and Jia Li. 2024. *The oscars of ai theater: A survey on role-playing with language models*.

Yang Deng, Wenqiang Lei, Minlie Huang, and Tat-Seng Chua. 2023a. *Goal awareness for conversational AI: Proactivity, non-collaborativity, and beyond*. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 6: Tutorial Abstracts)*, pages 1–10, Toronto, Canada. Association for Computational Linguistics.

Yang Deng, Wenqiang Lei, Minlie Huang, and Tat-Seng Chua. 2023b. *Rethinking conversational agents in the era of llms: Proactivity, non-collaborativity, and beyond*. In *Proceedings of the Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region, SIGIR-AP ’23*, page 298–301, New York, NY, USA. Association for Computing Machinery.

Yang Deng, Wenqiang Lei, Wai Lam, and Tat-Seng Chua. 2023c. A survey on proactive dialogue systems: Problems, methods, and prospects. *arXiv preprint arXiv:2305.02750*.

- Yang Deng, Wenqiang Lei, Lizi Liao, and Tat-Seng Chua. 2023d. [Prompting and evaluating large language models for proactive dialogues: Clarification, target-guided, and non-collaboration.](#)
- Yang Deng, Lizi Liao, Zhonghua Zheng, Grace Hui Yang, and Tat-Seng Chua. 2024. [Towards human-centered proactive conversational agents.](#) In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '24*, page 807–818, New York, NY, USA. Association for Computing Machinery.
- Yang Deng, Wenxuan Zhang, Wai Lam, See-Kiong Ng, and Tat-Seng Chua. 2023e. [Plug-and-play policy planner for large language model powered dialogue agents.](#) *arXiv preprint arXiv:2311.00262*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Esin Durmus, Karina Nyugen, Thomas I Liao, Nicholas Schiefer, Amanda Askill, Anton Bakhtin, Carol Chen, Zac Hatfield-Dodds, Danny Hernandez, Nicholas Joseph, et al. 2023. [Towards measuring the representation of subjective global opinions in language models.](#) *arXiv preprint arXiv:2306.16388*.
- Ritam Dutt, Sayan Sinha, Rishabh Joshi, Surya Shekhar Chakraborty, Meredith Riggs, Xinru Yan, Hao-gang Bao, and Carolyn Penstein Rosé. 2021. [Resper: Computationally modelling resisting strategies in persuasive conversations.](#) *arXiv preprint arXiv:2101.10545*.
- Marieke L Fransen, Edith G Smit, and Peeter WJ Verlegh. 2015. Strategies and motives for resistance to persuasion: An integrative framework. *Frontiers in psychology*, 6:1201.
- Yao Fu, Hao Peng, Tushar Khot, and Mirella Lapata. 2023. [Improving language model negotiation with self-play and in-context learning from ai feedback.](#)
- Lewis R Goldberg. 1992. The development of markers for the big-five factor structure. *Psychological assessment*, 4(1):26.
- He He, Derek Chen, Anusha Balakrishnan, and Percy Liang. 2018. [Decoupling strategy and generation in negotiation dialogues.](#) *arXiv preprint arXiv:1808.09637*.
- Zhiyuan Hu, Yue Feng, Yang Deng, Zekun Li, See-Kiong Ng, Anh Tuan Luu, and Bryan Hooi. 2023. [Enhancing large language model induced task-oriented dialogue systems through look-forward motivated goals.](#)
- Chen Huang, Peixin Qin, Yang Deng, Wenqiang Lei, Jiancheng Lv, and Tat-Seng Chua. 2024. [Concept – an evaluation protocol on conversational recommender systems with system-centric and user-centric factors.](#)
- Chen Huang, Peixin Qin, Wenqiang Lei, and Jiancheng Lv. 2023. [Reduce human labor on evaluating conversational information retrieval system: A human-machine collaboration approach.](#) In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 10876–10891, Singapore. Association for Computational Linguistics.
- Youngsoo Jang, Jongmin Lee, and Kee-Eung Kim. 2020. [Bayes-adaptive monte-carlo planning and learning for goal-oriented dialogues.](#) In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 7994–8001.
- Guangyuan Jiang, Manjie Xu, Song-Chun Zhu, Wenjuan Han, Chi Zhang, and Yixin Zhu. 2024. [Evaluating and inducing personality in pre-trained language models.](#) *Advances in Neural Information Processing Systems*, 36.
- Hang Jiang, Xiajie Zhang, Xubo Cao, Jad Kabbara, and Deb Roy. 2023. [Personallm: Investigating the ability of gpt-3.5 to express personality traits and gender differences.](#) *arXiv preprint arXiv:2305.02547*.
- Simon Keizer, Markus Guhe, Heriberto Cuayáhuilitl, Ioannis Efstathiou, Klaus-Peter Engelbrecht, Mihai Dobre, Alex Lascarides, and Oliver Lemon. 2017. [Evaluating persuasion strategies and deep reinforcement learning methods for negotiation dialogue agents.](#) In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 480–484, Valencia, Spain. Association for Computational Linguistics.
- Deuksin Kwon, Emily Weiss, Tara Kulshrestha, Kushal Chawla, Gale M Lucas, and Jonathan Gratch. 2024. [Are llms effective negotiators? systematic evaluation of the multifaceted capabilities of llms in negotiation dialogues.](#) *arXiv preprint arXiv:2402.13550*.
- Wenqiang Lei, Yao Zhang, Feifan Song, Hongru Liang, Jiabin Mao, Jiancheng Lv, Zhenglu Yang, and Tat-Seng Chua. 2022. [Interacting with non-cooperative user: A new paradigm for proactive dialogue policy.](#)
- Mike Lewis, Denis Yarats, Yann Dauphin, Devi Parikh, and Dhruv Batra. 2017. [Deal or no deal? end-to-end learning of negotiation dialogues.](#) In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2443–2453.
- Yu Li, Josh Arnold, Feifan Yan, Weiyan Shi, and Zhou Yu. 2021. [Legoeval: An open-source toolkit for dialogue system evaluation via crowdsourcing.](#) *arXiv preprint arXiv:2105.01992*.
- Yajiao Liu, Xin Jiang, Yichun Yin, Yasheng Wang, Fei Mi, Qun Liu, Xiang Wan, and Benyou Wang. 2023. [One cannot stand for everyone! leveraging multiple user simulators to train task-oriented dialogue systems.](#) In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1–21.

- Shima Rahimi Moghaddam and Christopher J Honey. 2023. Boosting theory-of-mind performance in large language models via prompting. *arXiv preprint arXiv:2304.11490*.
- OpenAI. 2022. Introducing chatgpt. <https://openai.com/blog/chatgpt>.
- OpenAI. 2023. Gpt-4 technical report. *ArXiv*, abs/2303.08774.
- David Premack and Guy Woodruff. 1978. Does the chimpanzee have a theory of mind? *Behavioral and brain sciences*, 1(4):515–526.
- Karsten Roth, Biagio Brattoli, and Bjorn Ommer. 2019. Mic: Mining interclass characteristics for improved metric learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*.
- Mustafa Safdari, Greg Serapio-García, Clément Crepy, Stephen Fitz, Peter Romero, Luning Sun, Marwa Abdulhai, Aleksandra Faust, and Maja Matarić. 2023. Personality traits in large language models. *arXiv preprint arXiv:2307.00184*.
- Maarten Sap, Ronan LeBras, Daniel Fried, and Yejin Choi. 2022. Neural theory-of-mind? on the limits of social intelligence in large lms. *arXiv preprint arXiv:2210.13312*.
- Susanne G Scott and Reginald A Bruce. 1995. Decision-making style: The development and assessment of a new measure. *Educational and psychological measurement*, 55(5):818–831.
- Weiyang Shi, Kun Qian, Xuewei Wang, and Zhou Yu. 2019. How to build user simulators to train rl-based dialog systems. *arXiv preprint arXiv:1909.01388*.
- Youzhi Tian, Weiyang Shi, Chen Li, and Zhou Yu. 2020. Understanding user resistance strategies in persuasive conversations. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4794–4798, Online. Association for Computational Linguistics.
- Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-sne. *Journal of machine learning research*, 9(11).
- Xintao Wang, Yaying Fei, Ziang Leng, and Cheng Li. 2023. Does role-playing chatbots capture the character personalities? assessing personality traits for role-playing chatbots. *arXiv preprint arXiv:2310.17976*.
- Xuewei Wang, Weiyang Shi, Richard Kim, Yoojung Oh, Sijia Yang, Jingwen Zhang, and Zhou Yu. 2019. Persuasion for good: Towards a personalized persuasive dialogue system for social good. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5635–5649, Florence, Italy. Association for Computational Linguistics.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2022. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*.
- Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8:229–256.
- Heinz Wimmer and Josef Perner. 1983. Beliefs about beliefs: Representation and constraining function of wrong beliefs in young children’s understanding of deception. *Cognition*, 13(1):103–128.
- Zelai Xu, Chao Yu, Fei Fang, Yu Wang, and Yi Wu. 2023. Language agents with reinforcement learning for strategic play in the werewolf game. *arXiv preprint arXiv:2310.18940*.
- Runzhe Yang, Jingxiao Chen, and Karthik Narasimhan. 2021. Improving dialog systems for negotiation with personality modeling.
- Xiao Yu, Maximillian Chen, and Zhou Yu. 2023. Prompt-based Monte-Carlo tree search for goal-oriented dialogue policy planning. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 7101–7125, Singapore. Association for Computational Linguistics.
- Haolan Zhan, Yufei Wang, Tao Feng, Yuncheng Hua, Suraj Sharma, Zhuang Li, Lizhen Qu, Zhaleh Semnani Azad, Ingrid Zukerman, and Gholamreza Haf-fari. 2024. Let’s negotiate! a survey of negotiation dialogue systems. *arXiv preprint arXiv:2402.01097*.
- Qiang Zhang, Jason Naradowsky, and Yusuke Miyao. 2023a. Ask an expert: Leveraging language models to improve strategic reasoning in goal-oriented dialogue models. *arXiv preprint arXiv:2305.17878*.
- Tong Zhang, Junhong Liu, Chen Huang, Jia Liu, Hongru Liang, Zujie Wen, and Wenqiang Lei. 2023b. Towards effective automatic debt collection with persona awareness. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: Industry Track*, pages 32–45, Singapore. Association for Computational Linguistics.
- Tong Zhang, Peixin Qin, Yang Deng, Chen Huang, Wenqiang Lei, Junhong Liu, Dingnan Jin, Hongru Liang, and Tat-Seng Chua. 2024. CLAMBER: A benchmark of identifying and clarifying ambiguous information needs in large language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10746–10766, Bangkok, Thailand. Association for Computational Linguistics.
- Xijia Zhang, Yue Guo, Simon Stepputtis, Katia Sycara, and Joseph Campbell. 2023c. Explaining agent behavior with large language models. *arXiv preprint arXiv:2309.10346*.

Pei Zhou, Aman Madaan, Srividya Pranavi Potharaju, Aditya Gupta, Kevin R McKee, Ari Holtzman, Jay Pujara, Xiang Ren, Swaroop Mishra, Aida Nematzadeh, et al. 2023a. How far are large language models from agents with theory-of-mind? *arXiv preprint arXiv:2310.03051*.

Xuhui Zhou, Hao Zhu, Leena Mathur, Ruohong Zhang, Haofei Yu, Zhengyang Qi, Louis-Philippe Morency, Yonatan Bisk, Daniel Fried, Graham Neubig, et al. 2023b. Sotopia: Interactive evaluation for social intelligence in language agents. *arXiv preprint arXiv:2310.11667*.

Yiheng Zhou, Yulia Tsvetkov, Alan W Black, and Zhou Yu. 2019. *Augmenting non-collaborative dialog systems with explicit semantic and strategic dialog history*.

A Details about Evaluation Protocol

A.1 Building User Simulators

Due to the significant human labor required for real-user evaluations (Huang et al., 2023), our experiments utilize user simulators instead.

A.1.1 Persona Generation

We prompt GPT4 (OpenAI, 2023) to generate diverse user personas by selecting attributes from two persona types, namely *Big-Five Personality* and *Decision-Making Styles*. Specifically, We allow GPT-4 to choose an attribute for each persona type, resulting in attribute-based user personas comprised of two fields, each containing a distinct attribute value. The prompt we use is provided in Table 11. In total, we create 20 attribute-based user personas and ensure that the number of each attribute is balanced. We then prompt GPT4 to rephrase these attribute-based personas into 300 cohesive persona descriptions. The prompt we use is provided in Table 12.

A.1.2 Non-collaborative Behavior Prompting

We leverage the *resisting strategies* outlined in (Dutt et al., 2021) as users’ non-collaborative behaviors. We provide the detailed explanations of these resisting strategies in Table 7. We design detailed instructions and incorporate these *resisting strategies* with their explanations into our user simulator prompting.

A.1.3 Comprehensive Prompting

By incorporating the persona description and resisting strategies, we construct comprehensive prompts for our user simulators. Specifically, our prompt includes two parts: task background and conversation history. In the task background, we guide LLMs

to role-play their assigned personas with a set of role-play instructions and resisting strategies. We provide the comprehensive user simulator prompts across two tasks in Table 13 and 14.

A.2 Evaluation Tasks

Following (Bianchi et al., 2024; Deng et al., 2023e), we consider two classic tasks as our evaluation scenarios, including price negotiation (He et al., 2018) and charity persuasion (Wang et al., 2019). The price negotiation task involves open-ended price negotiations where a buyer influences the seller towards a reasonable price, while the seller aims to maximize their own profit. The charity persuasion task involves asymmetric interactions guided by a persuader who endeavors to persuade the other party to make a charitable donation. Our evaluation is based on these two tasks, requiring the evaluated dialogue agents to take on the roles of buyer and persuader, respectively, in order to achieve their goals. To support our evaluations, we adopt the test dataset of CraigslistBargain (He et al., 2018) and PersuasionForGood (Wang et al., 2019), making use of their pre-annotated background information to streamline our assessment process. For the negotiation task, the background information includes item details and the desired price of each party. For the persuasion task, it involves determining if the individual being persuaded initially intends to make a donation. These background information serve as specific scenarios for our evaluation.

CB	Seller (User)	Buyer (Agent)
Target prices	285\$	142\$
Item	A skillfully lugged and elegantly pantographed road bike	
Goals	Maximize the price	Minimize the price
Ending condition	When either party accepts	
Max. # of turns	10 rounds of interaction	

Table 5: The evaluation scenario of price negotiation. This case is selected from the validate set of Craigslist-Bargain Dataset (He et al., 2018).

P4G	Persuader (Agent)	Persuadee (User)
Charity info	It works to help fight poverty around the world	
Goals	Convince the persuadee to donate	Foil the persuasion
Ending condition	When the persuadee agree to donate.	
Max. # of turns	10 rounds of interaction	

Table 6: The evaluation scenario of charity persuasion.

Resisting Strategy	Persuasion (P4G)	Negotiation (CB)
Source Derogation	Attacks/doubts the organisation’s credibility.	Attacks the other party or questions the item.
Counter Argument	Argues that the responsibility of donation is not on them or refutes a previous statement.	Provides a non-personal argument/factual response to refute a previous claim or to justify a new claim.
Personal Choice	Attempts to saves face by asserting their personal preference such as their choice of charity and their choice of donation.	Provides a personal reason for disagreeing with the current situation or chooses to agree with the situation provided some specific condition is met.
Information Inquiry	Ask for factual information about the organisation for clarification or as an attempt to stall.	Requests for clarification or asks additional information about the item or situation.
Self Pity	Provides a self-centred reason for not being able/willing to donate at the moment.	Provides a reason (meant to elicit sympathy) for disagreeing with the current terms.
Hesitance	Attempts to stall the conversation by either stating they would donate later or is currently unsure about donating.	Stalls for time and is hesitant to commit; specifically, they seek to further the conversation and provide a chance for the other party to make a better offer.
Self-assertion	Explicitly refuses to donate without even providing a factual/personal reason.	Asserts a new claim or refutes a previous claim with an air of finality/ confidence.
Others	Do not explicitly foil the persuasion attempts.	Do not explicitly foil the negotiation attempts.

Table 7: The resisting strategies for P4G and CB tasks.

Setting	Single-turn		Multi-turn	
	Natural	Useful	Natural	Useful
Human	18%	20%	15%	22%
TRIP	45%	42%	34%	31%
Tie	37%	38%	51%	48%

Table 8: Comparison on user simulators and real users. The Cohen’s Kappa between annotators is 0.67.

A.3 Reliability Analysis

Prior to conducting the interactive evaluation, we validate the reliability of using LLMs as user simulators that demonstrate non-collaborative behaviors. Following the approach described in (Deng et al., 2023e), we engage 5 human experts in conversations with two groups, including our diverse user simulators and 10 real users across two evaluation tasks. We collect 50 dialogues from each group and evaluate the user responses in both single-turn and multi-turn open-ended conversations. The evaluation focuses on the naturalness and utility of the generated responses in these conversation settings. Naturalness refers to the fluency and human-like nature of the responses, while utility indicates their consistency with the role instructions and non-collaborative behaviors. We employ two annotators to conduct pairwise evaluations by rating "Win/Tie/Lose" between the two samples. As shown in Table 8, the user simulators exhibit a notably superior performance compared to real users, particularly when it comes to the naturalness of responses in multi-turn conversations, which showcases the impressive language generation ca-

pabilities inherent in LLMs. Furthermore, even compared with human-annotated dialogues, the GPT3.5-based simulator shows competitive performance. These results validate the reliability of adopting GPT3.5 as the user simulator.

A.4 Interactive Evaluation Protocol

During the evaluation, each dialogue agent must engage with these simulators (Deng et al., 2023e). During interactions, the dialogue agent and user simulator alternate in employing strategies in their responses with the ultimate aim of maximizing their own self-interest. The interactions continues until the conversational goal is achieved or the maximum number of turns T (i.e., T is set to 10 for both tasks) is reached. To determine goal achievement, we utilize AI feedback to assess whether the task goal has been reached. Specifically, in price negotiation task, we employ a separate GPT3.5 (i.e., LLM_{rwd}) to assess whether both parties have reached a deal. We prompt LLM_{rwd} to generate feedback for the binary question "Have they reached a deal?". If the output of LLM_{rwd} indicates that both parties have reached an agreement, we consider this as goal achievement. In charity persuasion task, we additionally prompt the user simulator to express his willingness to make a donation at the end of each turn. In particular, we query the user simulator "Would you be interested in donating to Save the Children?". If the feedback is positive, we regard this as goal achievement. Conversely, if the goal is not achieved, the interaction continues.

Due to the subjectivity of the planning outcome as well as the variance of the LLM-generated output, we follow a common practice (Wang et al., 2022; Deng et al., 2023e) to alleviate these issues by sampling the decoded sequences l (i.e., l is set to 10 for both tasks) times.

B Implementation Details

B.1 TRIP Implementation Details

B.1.1 Theory-of-Mind

We leverage the strong Theory-of-Mind capability of GPT3.5 to infer the mental states and user future actions during interaction. The prompt we use is provided in Table 15 and 16.

B.1.2 Strategy Prompting

Here, we present the dialogue agent strategies utilized in our experiments. Initially, we outline the strategies along with their explanations for two tasks in Table 9 and 10. We then offer a comprehensive overview of our TRIP prompting in Table 19 and 20.

B.1.3 Supervised Fine-Tuning

We initialize our strategy planner by imitating human-human dialogue datasets in CraigslistBargain and PersuasionForGood through supervised fine-tuning (SFT). In specific, we adopt the strategy annotations in the train dataset to support our SFT. We optimize the strategy planner by minimizing the cross-entropy loss between the predicted strategy y_i and the human annotated strategy \hat{y}_i :

$$\mathcal{L}_{CE} = -\frac{1}{m} \sum_{i=1}^m [y \log \hat{y}_i + (1 - y_i) \log(1 - \hat{y}_i)]$$

Regarding the training hyper-parameters, we set the batch size 16 and the learning rate 6e-6, and utilize the AdamW optimizer with a weight decay of 0.01. We save the checkpoint based on the best performance at the validation set.

B.1.4 Online RL Training

After SFT, we optimize our strategy planner through REINFORCE algorithm. In specific, our training involves 1000 episodes, with a learning rate of 1e-6, a discount factor 0.999, and the maximum conversation turn of each episode 10. All the training experiments are run on a server equipped with 4 Tesla V100 GPUs.

B.2 Baselines Implementation Details

We implement the existing LLM-based dialogue agents by following previous works.

Standard: simply prompts LLMs to chat with users using task instructions without considering any dialogue strategy.

ProCoT: we follow (Deng et al., 2023d) and prompt LLM to analyze the dialogue status and plan next strategy, and then generate a response based on the planned strategy. We provide its prompt design in Table 17.

ICL-AIF: we follow (Fu et al., 2023) and prompt another GPT3.5 for verbal feedback, offering suggestions to the dialogue agent upon completion of an interaction. Our implementation involves presenting three suggestions at the conclusion of each interaction, while ensuring that only the most recent 20 suggestions are retained to prevent indefinite expansion. The prompt we use is provided in Table 18.

GDP-MCTS: we follow (Yu et al., 2023) and implement open-MCTS to help LLM for strategic planning. This method is originally proposed for charity persuasion dialogues. In order to further accommodate the price negotiation applications, we just need to modify the task instruction and the role-playing description.

PPDPP: we follow (Deng et al., 2023e) and adopt the BERT⁷ model (Devlin et al., 2018) as our external planner. We implement PPDPP based on the training details provided in the original paper. We have made adjustments to the task instructions and role-playing descriptions, adapting them for use in the context of charity persuasion.

C Human Evaluation

Inspired by (Yu et al., 2023), we conduct interactive human evaluation using the LegoEval platform (Li et al., 2021) with crowdworkers on Amazon Mechanical Turk. We primarily sought to evaluate TRIP against two competitive baselines (i.e., *Standard* and *PPDPP*). In specific, we hire 20 crowdworkers with varying personas to converse with our three agents based on the price negotiation and charity persuasion tasks. After conversations, we collect 50 dialogues for each agent and calculate their performances using the same metrics mentioned in Section 3.1.

⁷<https://huggingface.co/google-bert/bert-base-uncased>

D More Experimental Results

In addition to the Success Rate, we report the agents performance across various personas using the metrics of Average Turn and Sale-to-List Ratio, as depicted in Figure 8 and Figure 7. We discover that the overall performance and analysis conclusions remain largely consistent with Section 5.1.

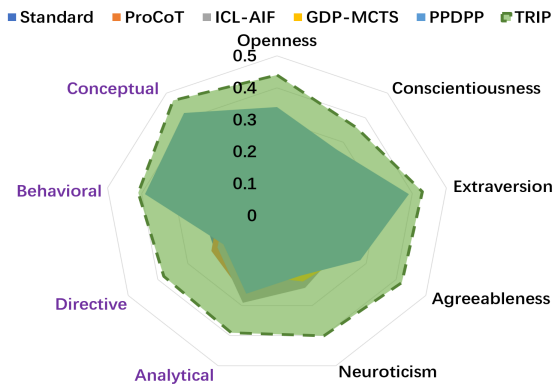


Figure 7: The agents performance across various personas. We report their SL % on the price negotiation task. TRIP achieves balanced improvements on all personas, significantly outperforming other agents by a considerable margin.

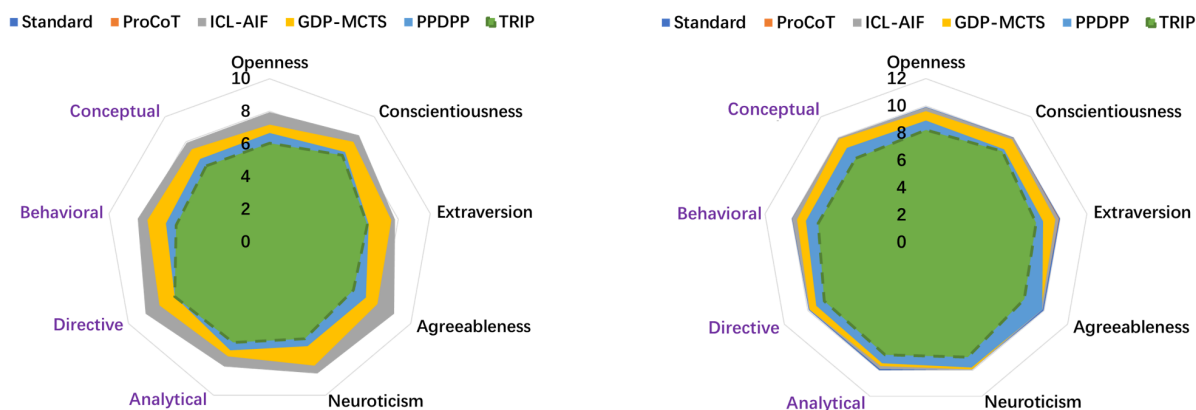


Figure 8: The agents performance across various personas. We report their average turn on two tasks, namely price negotiation (*Left*) and charity persuasion (*Right*). TRIP achieves balanced improvements on all personas, significantly outperforming other agents by a considerable margin.

Dialogue Strategy	Explanation
Greetings	Please say hello or chat randomly.
Ask a question	Please ask any question about product, year, price, usage, etc.
Answer a question	Please provide information about the product, year, usage, etc.
Propose the first price	Please initiate a price or a price range for the product.
Propose a counter price	Please propose a new price or a new price range.
Use comparatives	Please propose a vague price by using comparatives with existing price.
Confirm information	Please ask a question about the information to be confirmed.
Affirm confirmation	Please give an affirmative response to a confirm.
Deny confirmation	Please give a negative response to a confirm.
Agree with the proposal	Please agree with the proposed price.
Disagree with a proposal	Please disagree with the proposed price.

Table 9: The negotiation strategies used in our TRIP agent.

Dialogue Strategy	Explanation
Logical Appeal	Please use of reasoning and evidence to convince the persuadee.
Emotion Appeal	Please elicit the specific emotions to influence the persuadee.
Credibility Appeal	Please use credentials and cite organizational impacts to establish credibility and earn the user's trust. The information usually comes from an objective source (e.g., the organization's website or other well-established websites).
Foot in the Door	Please use the strategy of starting with small donation requests to facilitate compliance followed by larger requests.
Self-Modeling	Please use the self-modeling strategy where you first indicates the persuadee own intention to donate and chooses to act as a role model for the persuadee to follow.
Personal Story	Please use narrative exemplars to illustrate someone donation experiences or the beneficiaries positive outcomes, which can motivate others to follow the actions.
Donation Information	Please provide specific information about the donation task, such as the donation procedure, donation range, etc. By providing detailed action guidance, this strategy can enhance the persuadee's self-efficacy and facilitates behavior compliance.
Source-related Inquiry	Please ask if the persuadee is aware of the organization (i.e., the source in our specific donation task).
Task-related Inquiry	Please ask about the persuadee opinion and expectation related to the task, such as their interests in knowing more about the organization.
Personal-related Inquiry	Please asks about the persuadee previous personal experiences relevant to charity donation.

Table 10: The persuasion strategies used in our TRIP agent.

The prompt for user persona generation

You need to select one attribute from each of the following persona types.

Persona types

Big-Five Personality: ["openness", "conscientiousness", "extraversion", "agreeableness", "neuroticism"]

Decision-Making Styles: ["directive", "analytical", "conceptual", "behavioral"]

Please generate a list of N fictional user profiles.

Table 11: The prompt of user persona generation.

The prompt for user persona rephrase

You need to incorporate the following persona attributes and generate a cohesive persona description. You need to ensure the description is easy to understand.

Big-Five Personality:

Decision-Making Style:

An Example:

You are a 28-year-old female software developer. Your personality is characterized by openness to experience, which means you are curious, imaginative, and willing to try new things. In your occupation, you excel at analyzing problems and finding logical solutions. Your decision-making style is analytical, meaning you carefully consider all available information before making a choice.

Table 12: The prompt of user persona rephrase.

The user simulator prompt for the price bargain task

Now enter the role-playing mode. In the following conversation, you will play as a seller in a price bargaining game.

Your persona: <Persona Description>

You must follow the instructions below during chat.

1. Your utterances and bargain behavior need to strictly follow your persona. Varying your wording and avoid repeating yourself verbatim!
2. You can decide to change your target price flexibly based on your persona and the conversation.

Your Response Strategy:

1. "Source Derogation": Attacks the other party or questions the item.
2. "Counter Argument": Provides a non-personal argument/factual response to refute a previous claim or to justify a new claim.
3. "Personal Choice": Provides a personal reason for disagreeing with the current situation or chooses to agree with the situation provided some specific condition is met.
4. "Information Inquiry": Requests for clarification or asks additional information about the item or situation.
5. "Self Pity": Provides a reason (meant to elicit sympathy) for disagreeing with the current terms.
6. "Hesitance": Stalls for time and is hesitant to commit; specifically, they seek to further the conversation and provide a chance for the other party to make a better offer
7. "Self-assertion": Asserts a new claim or refutes a previous claim with an air of finality/ confidence.
8. "Others": Do not explicitly foil the negotiation attempts.

You are the seller who is trying to sell the %s with the initial price of %s. Product description: %s. Please reply with only one short and succinct sentence.

Conversation History

Table 13: The comprehensive prompt of user simulators in the price negotiation task.

The user simulator prompt for the charity persuasion task

Now enter the role-playing mode. In the following conversation, you will play as a Persuadee in a persuasion game.

Your persona: <Persona Description>

You must follow the instructions below during chat.

1. Your utterances need to strictly follow your persona. Varying your wording and avoid repeating yourself verbatim!
2. Pretend you have little knowledge about the Save the Children charity. You have little willingness for donation at the beginning of conversation.
3. Your willingness for donation depends on your persona and how influenced you are by the Persuader.
4. You need to determine whether to donate money or not. If you think you are convinced by the Persuader, you should donate money for the charity.

Your Response Strategy:

1. "Donate": show your willingness to donate.
2. "Source Derogation": attacks or doubts the organisation's credibility.
3. "Counter Argument": argues that the responsibility is not on them or refutes a previous statement.
4. "Personal Choice": Attempts to save face by asserting their personal preference such as their choice of charity and their choice of donation.
5. "Information Inquiry": Ask for factual information about the organisation for clarification or as an attempt to stall.
6. "Self Pity": Provides a self-centred reason for not being willing to donate at the moment.
7. "Hesitance": Attempts to stall the conversation by either stating they would donate later or is currently unsure about donating.
8. "Self-assertion": Explicitly refuses to donate without even providing a personal reason.
9. "Others": Do not explicitly foil the persuasion attempts.

You are the Persuadee who is being persuaded by a Persuader. Please reply with only one short and succinct sentence.

Conversation History

Table 14: The comprehensive user simulator prompt for the charity persuasion task.

The Theory-of-Mind prompt for the price negotiation task

You are an expert in price bargain.

Now give you a conversation history between a buyer and a seller, you need to infer the mental states and future actions of the seller.

Conversation History

Table 15: The ToM prompt for the price negotiation task.

The Theory-of-Mind prompt for the charity persuasion task

You are an expert in charity persuasion.

Now give you a conversation history between a persuader and a persuadee, you need to infer the mental states and future actions of the persuadee.

Conversation History

Table 16: The ToM prompt for the charity persuasion task.

The prompt of the ProCoT agent

The Price Negotiation Task

Assume you are the buyer. Given the conversation history, in order to reach a better deal with the seller, please select the most appropriate dialogue strategy.

You can only reply by selecting one of the following dialogue strategy to reach the goal: Greetings. Ask a question. Answer a question. Propose the first price. Propose a counter price. Use comparatives. Confirm information. Affirm confirmation. Deny confirmation. Agree with the proposal. Disagree with a proposal.

The following is the conversation history: [conversation]

The Charity Persuasion Task

Assume you are the Persuader. Given the conversation history, in order to convince the persuadee to donate for charity, please select the most appropriate dialogue strategy.

You can only reply by selecting one of the following dialogue strategy to reach the goal: Logical appeal, Emotion appeal, Credibility appeal, Foot in the door, Self-modeling, Personal story, Donation information, Source-related inquiry, Task-related inquiry, Personal-related inquiry.

The following is the conversation history: [conversation]

Table 17: The prompt design of the *ProCoT* agent.

The prompt of the ICL-AIF agent

The Price Negotiation Task

Now enter the role-playing mode. In the following conversation, you will play as a coach in a bargain game. There will be a buyer and a seller bargaining about a product price.

Your task is to read the conversation between the buyer and the seller, then provide suggestions to the buyer about how to buy the product with a lower price. Each suggestion should be only one short and succinct sentence.

The following is the conversation: [conversation]

Question: What are your suggestions? Answer:

The Charity Persuasion Task

Now enter the role-playing mode. In the following conversation, you will play as a coach in a persuasion game. There will be a persuader who is trying to persuade a persuadee for charity donation. Your task is to read the conversation between the persuader and the persuadee, then provide suggestions to the persuader about how to convince the persuadee to make a donation. Each suggestion should be only one short and succinct sentence.

The following is the conversation: [conversation]

Question: What are your suggestions? Answer:

Table 18: The prompt design of the *ICL-AIF* agent.

The prompt of our TRIP agent in price negotiation

Now enter the role-playing mode. In the following conversation, you will play as a buyer in a price bargaining game.

You are the buyer who is trying to buy the %s with the price of %s. Product description: %s

Please reply with only one short and succinct sentence. [action] Now start the game.

Table 19: The prompt design of the TRIP agent for price negotiation.

The prompt of our TRIP agent in charity persuasion

Now enter the role-playing mode. In the following conversation, you will play as a Persuader who is trying to persuade the Persuadee to donate to the charity called Save the Children.

Save the Children is head-quartered in London, and they work to help fight poverty around the world. Children need help in developing countries and war zones. Small donations like \$1 or \$2 go a long way to help.

You are the Persuader who is trying to convince the Persuadee to donate to a charity called Save the Children. [action]

Please reply with only one short and persuasive sentence.

Table 20: The prompt design of the TRIP agent for charity persuasion.